Soc Choice Welfare (1988) 5:45–68

Social Choice and Welfare

© Springer-Verlag 1988

Strategic Voting in Repeated Referenda

S. Morton*

Department of Accounting and MIS, The Pennsylvania State University, College of Business Administration, University Park, PA 16802, USA

Received March 29, 1985 / Accepted October 12, 1987

Abstract. This paper analyzes learning and voting strategy when a budget maximizing bureaucrat has several chances to obtain referendum approval. The process is modeled as a sequential game with a continuum of heterogeneous voters and a dominant bureaucrat in which all agents are uncertain about the true distribution of voter preferences. The equilibrium concept is perfect Bayes Nash, so voting is strategic in the sense of foresighted but nevertheless noncooperative.

1. Introduction

Public opinion polls are increasingly common in larger federal and state elections. In addition to their entertainment value, polls play an important role in political strategy since the expense of the polls is often borne by the candidate or party. This indicates that there is a significant element of uncertainty about voter preferences among political strategists. This essay begins to analyze the interplay of information acquisition and strategy in a political process by constructing a Bayes equilibrium of repeated referenda in which agents acquire information on voter preferences by observing the results of early referenda and use that information in formulating strategy for subsequent referenda.

Recently, Romer and Rosenthal [11] have introduced a monopoly model of a voting process which is better suited to this analysis than the earlier competitive nature of median voter theories, since competitive theories tend to abstract away from strategic behavior. Romer and Rosenthal analyzed a referendum process in which a budget maximizing bureaucrat must win majority approval from the voters for the level of public expenditure. The ideal point of the median voter is the only alternative that can defeat any other, so if there is free entry in placing proposals on

* Thanks to R. Gretlein, J. Hamilton, T. Palfrey, T. Romer, H. Rosenthal, S. Slutsky and the referees of this journal for their comments and encouragement. They are not responsible for any errors or omissions.

the ballot, the tastes of the median voter should prevail. In many instances, however, the bureaucrat has agenda control, and can prevent some alternatives from being proposed that would have defeated his favorite. Romer and Rosenthal showed that a bureaucrat will gain some monopoloy power with agenda control and, depending on the distribution of voter tastes, will be able to obtain a considerably larger budget than the median voter's ideal.

This essay analyzes learning and strategy in the political process by considering repeated referenda in which the budget maximizing agenda setter has two chances to obtain majority approval. Voting terminates with the first proposal that succeeds or a prespecified reversion level if both fail, assumed here to be normalized to zero. If agents had perfect knowledge about the distribution of voter tastes then one would never expect to see the first vote fail since the setter would propose on the first round the highest proposal he knows will pass with 50% of the vote. Therefore it will be assumed that all agents are unsure of the distribution. Learning occurs since agents will observe the vote for the first proposal and make inferences about the distributions that would have generated this vote, given equilibrium voting strategies. The agenda setter uses this information to formulate strategy for the second round of voting if the first proposal failed. The voters are also strategic since they take account of this signalling function of the first vote in formulating their voting strategies.

Romer and Rosenthal analyzed repeated referenda but abstracted away from both learning and strategy in the sense analyzed here. They assumed voters were myopic or naive and voted in each round as if it were the last, and that the setter did not learn from the previous vote. When agents are more fully rational, both behaviors are more complex. In the last round, if there is one, voting is myopic since voters make a straightforward comparison between the current proposal and reversion. However, the second proposal will be chosen on the basis of the setter's (Bayesian) updated information about the distribution of tastes, in contrast to Romer and Rosenthal for whom the last proposal is identical to the one chosen if there had been only one round of voting, since they assume no learning.

In the first round, voting is far from myopic since foresighted voters must make some conjecture about what will occur in the second round in order to compare it to the first proposal. This conjecture will be a "rational expectation" in which voters must predict what proposal the setter will make conditional on the information gained from all the possible first period votes. In fact, voters may cast a vote in order to influence the signal to the setter and the choice of the second proposal. This raises the possibility of a "protest" vote in which voters object to the first proposal which they would myopically prefer to reversion, in order to obtain or manipulate a better second proposal. At the same time, the first vote is not a pure public opinion poll and voters must also consider that the first proposal will pass. They must weigh the trade off between voting for a proposal in order to obtain its passage versus influencing the choice of the next proposal.

The well known problem of why voters vote introduces additional complexity into this analysis of strategy and learning. Both Romer and Rosenthal and this essay assume a continuum of voters for the sake of tractability of the analysis. The issue of why voters vote when each individually has no weight in the outcome is finessed in the usual way by assuming there are no costs of voting and perhaps psychic

benefits[1]. Rational voting is motivated either as voting one's preferences or voting as if one were decisive. When there are two determinate alternatives these two heuristics are equivalent; for repeated referenda, however, the alternative to the first proposal is a lottery over reversion and the second proposal (which may itself be random since it depends on the information communicated by the results of the first vote). Voting one's preferences entails comparing the value of the current proposal to the expected value of the lottery. On the other hand, the assumption that one will be decisive implies that all other voters are evenly split, and requires implicit assumptions about the otherwise unknown distribution of voter tastes. Thus, these two heuristics may not be equivalent here. Although the correct approach must await a theory of why voters vote, this essay uses an explication of rational voting in the first period suggested by a limiting analysis of finite voting populations where voters may have weight in the outcome.

The two-stage voting process is modeled here as a sequential game with a continuum of heterogeneous voters and a dominant-player bureaucrat. All agents are uncertain about the true distribution of voter tastes and it is assumed that agents share a common prior about the unknown parameters of the distribution. To model learning, Bayes Rule is used by agents to update their prior after seeing the first vote. Also, it is assumed that agents do not or cannot commit themselves in advance to take some specific or contingent action at a later date (e.g., the setter "threatening" that if his first proposal is not accepted he will refuse to make a second proposal) so the equilibrium concept will be perfect. Finally, all agents behave noncooperatively, or Nash, which rules out coalitions or voting blocs from forming to decide on cooperative voting strategies which may not be individually rational. Thus, voting is strategic or sophisticated in the special sense of foresighted (i.e., with "rational expectations") but nevertheless noncooperative.

The perfect Bayes Nash equilibrium concept used here has been formalized by Kreps and Wilson [4], and employed in other contexts in which learning and strategy interact, especially in private information models of bargaining, such as Cramton [1], Fudenberg and Tirole [3], and Morton [8] and in some reputation models in industrial organization, such as Kreps and Wilson [5] and Milgrom and Roberts [7]. Although this essay seems to be the first to address these issues in a voting model the underlying structure of the model is similar and requires the solution approach common to that literature – solving the model backwards for every possible last period information set and then examining how that information set may be reached from prior period behavior. Only in this way can agents address the tradeoff in choosing an action to obtain a desirable future information set versus choosing an action for its immediate allocation value. This is critical to both reputation building and dynamic bargaining strategies, and is also of central concern here. In this light, it is quite interesting to note Crawford and Sobel [2], who analyze a "pure" information transmission problem, analogous here to a pure public opinion poll or "straw vote" in which the voters send a message to the setter by responding to a poll. The aggregated information received by the setter is then used to propose a budget. They found that if there is sufficient divergence in agents'

[1] See Ledyard [6] or Palfrey and Rosenthal [10] for game theoretic analyses of costly voting and strictly rational voting strategies.

interests then foresighted players will never volunteer any of their private information. This suggests that a pure public opinion poll of strategic voters will have no value.

The structure of the paper is as follows: the next section specifies the environment. Then the following section presents an introductory version of the model in which there is a minimal amount of initial uncertainty about the distribution of voter tastes so that observing the results of the first vote is sufficient to give agents perfect information in the second round of voting. Then Sect. 4 specifies the perfect Bayes equilibrium for the general case and addresses the issue of how rational voters should vote without any weight in the outcome. Section 5 constructs an equilibrium under some specific assumptions about the initial information structure and Sect. 6 sketches the proof of uniqueness. The final section concludes with a discussion of generalizations as well as empirical implications of strategic learning.

2. The Environment

There is a risk neutral setter or bureaucrat who wishes to maximize the size of the budget or expenditure level, e. In order to do so, the setter has two chances to win the approval of a majority of the voters. If majority approval is not obtained, then the budget level will be the reversion level, normalized to be zero expenditure.

There is a continuum of voters, each of which is characterized by a political preference parameter, t. A voter of type t has a quadratic utility function for expenditure $u(e, t) = te - e^2$. These are single-peaked preferences and have the characteristic that $t/2$ is the ideal expenditure for a voter of type t, and $e = t$ is the largest expenditure level which a voter of type t will prefer to the reversion budget of zero.

It is common knowledge among the voters and the setter that voter tastes, t, are distributed in the voting population according to some continuous p.d.f., $t \sim f(\cdot \mid y)$ where y is a finite vector of parameters. However, y is unknown and the agents' common beliefs are that y is distributed according to some joint p.d.f., $y \sim g(y)$. For example, it may be known that tastes are distributed normally in the population but the mean and the variance are unknown. Each voter observes his private preference parameter, t, but does not use this observation to update his beliefs about the distribution of tastes; that is, if a voter's t is extreme under the common knowledge distribution, then the voter understands that he is extreme.

The setter has at most two opportunities to hold referenda on proposed expenditure levels. If the setter's proposed budget, e_1, passes in the first period with a majority of the vote, then that will be the budget throughout the duration of the budget life, which is assumed to be fixed exogenously. If the first period proposal fails, then the setter proposes a second budget, e_2, which, if passed, will be the budget level for the remainder of the budget life. If neither passes, the budget will be the reversion level of zero[2].

[2] It is straightforward, but unenlightening, to impose election costs on the setter or discounting on the voters. Discounting may be appropriate if, for example, the public good or service is not supplied during the interim between a failed first referendum and the second, in which case it represents the opportunity cost of the interrupted flow of public good or service which is funded by the budget.

Finally, it is assumed that all voters vote according to the state they prefer. Since a solution to the problem of why voters vote is not offered here, a referendum is to be thought of as a costless and exhaustive, but strategic, public opinion poll. In the last period, voting is simply myopic as in Romer and Rosenthal, where voters behave as if their vote will be decisive. In the first period, preference comparisons are more subtle since there are more states to compare: the current proposal, the uncertain proposal of the next period or reversion. Unlike the last period, voters must also be aware of the strategic function of their vote in influencing the next proposal. Section 4 discusses how this preference comparison should be modeled.

3. A Case of Perfect Learning

As a prolegomenon to the more general analysis, this section considers the case in which agents have relatively little uncertainty from the outset, so that observation of the results of the first round of voting allows agents to infer the actual distribution of voter tastes. Although this simplified case is useful as an intermediate step between Romer and Rosenthal's myopic behavior and the imperfect learning considered later, it is not plausible that agents will have perfect information after only one round of voting; also it does not capture some of the more interesting features of learning.

To make foresighted calculations in the first period, agents must be able to forecast the future. So consider the second period first under the assumption that the distribution of voter tastes is known. Voters will behave myopically since there is no further future and will vote in favor of the second proposal, e_2, if they prefer it to reversion, or for the utility function assumed here, if $t \geqq e_2$. The second proposal will pass under majority rule if more than half of the voters vote for it or $F(e_2|y) \geqq 1/2$, where F is the cumulative of f, the distribution of voter tastes. This can be expressed as: the median of (the distribution of) t given y is greater than e_2.

Since the setter is presumed to know the distribution of t in the second round, he also knows precisely which proposals will pass. Seeking to maximize the budget, he will propose the largest budget that will pass, or $e_2 =$ the median of t, which will surely just pass.

In fact, the setter need not know the entire distribution, but only its median, in order to choose the second proposal. Presumably the first vote will have conveyed information about the median which was previously unknown. Therefore, it will be assumed in this section that the vector of parameters, y, of the distribution of tastes is just a single parameter, which is most conveniently interpreted as the median. The support of the prior distribution of the median is assumed to be non-negative. It is important for perfect learning that agents be uncertain of only one parameter of the distribution, so that its true value can be inferred from the single observation of the first vote. Additional assumptions necessary for this inference will be introduced below.

Turning to the first period, consider the voters' decision problem. The setter has proposed some e_1 and voters will vote their preferences by comparing e_1 to the expected value of what will occur in the future; specifically, the second proposal will be whatever the median of voter tastes is revealed to be and it will surely pass.

Therefore, the voters' decision in the first period is identical to a single stage referendum process in which reversion is the random variable $y =$ median of t. (This is the sense in which this case of perfect learning trivializes some of the strategic aspects of the voters' learning process.) A voter will favor the first proposal if $\Delta(e_1, t) \geqq 0$, where

$$\begin{aligned}\Delta(e_1, t) &\equiv u(e_1, t) - Eu(y, t) \\ &= (e_1 t - e_1^2) - E(yt - y^2) \\ &= (e_1 - Ey)t - (e_1^2 - Ey^2) \ .\end{aligned}$$

Of particular interest is the switching voter – that type, $t^*(e_1)$, who is indifferent between the current proposal and the random reversion level, y,

$$\Delta(e_1, t^*(e_1)) \equiv u(e_1, t^*(e_1)) - Eu(y, t^*(e_1)) \equiv 0 \ .$$

By this definition, the switching type views e_1 as the certainty equivalent of the lottery y; i.e., as equal to the mean, $\bar{y}$, plus an adjustment for risk. For the quadratic utility function assumed here, it is easily verified that this risk adjustment is just equal to the variance of y:

$$Eu(y, t) = u(\bar{y}, t) - \sigma^2$$

for any t and distribution of y, where σ^2 is the variance of y, since in general $\sigma^2 = Ey^2 - \bar{y}^2$. Thus, the switching type is defined by

$$u(e_1, t^*(e_1)) = u(\bar{y}, t^* e_1) - \sigma^2$$

or

$$\begin{aligned}t^*(e_1) &= (e_1^2 - \bar{y}^2 - \sigma^2)/(e_1 - \bar{y}) \\ &= e_1 + \bar{y} - \sigma^2/(e_1 - \bar{y}) \ .\end{aligned}$$

The graph of $t^*(e_1)$ is in Fig. 1. Since Δ is linear in t, there is at most one switching type for every first proposal. All types above $t^*(e_1)$ will vote together in favor of the first proposal if the coefficient on t in Δ is positive or $e_1 > \bar{y}$; conversely, if $e_1 < \bar{y}$. This voting strategy implies that the fraction of the vote in favor of the first proposal is

$$V_1 = \begin{cases} 1 - F(t^*(e_1)|y) & \text{if} \quad e_1 > \bar{y} \\ F(t^*(e_1)|y) & \text{if} \quad e_1 < \bar{y}^3 \ . \end{cases}$$

e_1 will pass if more voters are in favor than against, or equivalently, if the unknown median type, $t = y$, is in favor. So,

$$e_1 \text{ passes iff } \begin{cases} y \geqq t^*(e_1) & \text{if} \quad e_1 > \bar{y} \\ y \leqq t^*(e_1) & \text{if} \quad e_1 < \bar{y} \end{cases} .$$

[3] The case $e_1 = \bar{y}$ causes every voter to vote for e_1 since $\Delta(\bar{y}, t) > 0$ for all t. But the setter can do better, see below.

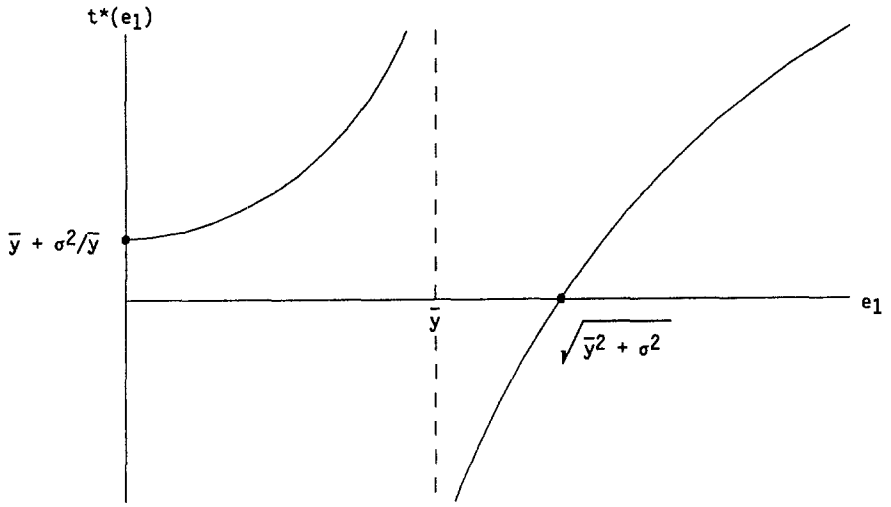

Fig. 1. The switching type for the case of perfect learning

Another assumption is required, to ensure that y can be inferred from V_1. Perhaps the most reasonable is that an increasing median causes F to (first order) stochastically increase: for all t

$$F(t|y) < F(t|y')$$

when $y < y'$, as in the normal distribution. Thus, if the coefficient of t in Δ is positive a given proposal will receive more votes as the distribution shifts to the right and the median taste increases; conversely if the coefficient is negative.

Since the setter chooses e_1, he can control the form of voting by choosing e_1 larger or smaller than $\bar{y}$. In every case, however, the information contained is exactly the same – the true median is revealed. It is clear that a budget maximizing setter will never choose $e_1 < \bar{y}$, since he can obtain the same $t^*(e_1)$ and probability of passage in the first period, and the same information for the future, by choosing some higher $e_1 > \bar{y}$.

In fact, the lowest proposal that the setter might choose is $e_1 = (\bar{y}^2 + \sigma^2)^{1/2} > \bar{y}$, for which $t^*(e_1) = 0$. Since it has been assumed that $y \geqq 0$, the true median will always be greater than this switching type above which voters are in favor, so this proposal will surely pass in the first round.

It is remarkable that the setter can be guaranteed that such a high first proposal will pass. The explanation is that the setter is able to use the second round as a threat against risk averse voters. All agents have a common information set and the expected outcome of the next period lottery is $\bar{y}$. Even the type $t = 0$ who prefers zero expenditure to any positive level will be willing to pay a risk premium of $(\bar{y}^2 + \sigma^2)^{1/2} - \bar{y}$ to avoid that lottery. Essentially the original reversion level has become irrelevant.

This first proposal may or may not solve the setter's problem. By choosing a higher e_1, the setter may win a larger budget in the first round, but the probability of failure also begins to increase and in the second round the risk neutral setter will

receive the lower $\bar{y}$, the expected value of the lottery. The setter's problem is to choose $e_1 \geqq (\bar{y}^2+\sigma^2)^{1/2}$ to

$$\max E_y \begin{cases} e_1 & \text{if } y \geqq t^*(e_1) \\ y & \text{if } y < t^*(e_1) \end{cases} = e_1[1-G(t^*(e_1))] + \int_0^{t^*(e_1)} y dG(y) \ ,$$

where $G(\cdot)$ is the cumulative of the distribution $g(y)$.

The first order conditions are

$$g(t^*(e_1))\{[1-G(t^*(e_1))]/g(t^*(e_1)) - [e_1 - t^*(e_1)]t^{*\prime}(e_1)\} = 0 \ .$$

These conditions bear a close resemblance to other private information models. The left hand term in the braces is the reciprocal of the hazard rate of the distribution $g(y)$ and the "regularity" or second order condition that it is strictly decreasing is commonly assumed. For example, the first order condition for a one stage referendum (see Romer and Rosenthal [11]) is simply

$$[1-G(e)]/g(e) - e = 0 \ .$$

The right hand term is strictly increasing so a unique solution is guaranteed[4]. Most of the common distributions are everywhere "regular". In the two stage referendum, however, the right hand term is also strictly decreasing, so even with regularity there may be multiple solutions to the first order conditions, each of which must be compared to the others and the corner, or even no solution if the support of the distribution $g(y)$ has no upper bound. As an example, if y has the uniform distribution, the first order condition becomes a quartic polynomial in e_1.

This example has introduced some of the features of the more general analysis, such as foresighted voting and proposing, and the existence of a switching type of voter. However, learning is implausibly easy in this model, which perhaps gives the setter too much control. Also, the distribution of the second proposal is purely exogenous which restricts the sense in which the model is strategic and weakens its empirical implications. The more general analysis follows.

4. Perfect Bayes Equilibria

It is easiest to specify the more general definition of equilibrium backwards. In the second and last period, there is no future so the voting strategy is myopic. Although a voter has no weight in the outcome, in the unlikely event of being decisive voters are sincere. A voter of type t votes in favor of the second proposal, e_2, if he prefers it to reversion:

$$v_2(e_2, t) = \begin{cases} 1 & \text{if } u(e_2, t) \geqq u(0, t) \quad \text{or} \quad t \geqq e_2 \\ 0 & \text{otherwise} \end{cases} \ .$$

Then e_2 will pass if and only if there are more voters with type greater than e_2 than less, or $F(e_2|y) < 1/2$, where F is the cumulative of f and is conditioned on the true population parameters, y.

[4] It is also used in a simple auction model in which a risk neutral seller chooses the price e to sell to buyers with unknown reservation values. (See Myerson [9] for this early use of regularity.)

A risk neutral setter will want to propose the largest expenditure level that will pass, but may still have some residual uncertainty of what this level is, because of uncertainty about y. (This is the most interesting case.) In the last period, the setter will choose e_2 to solve

$$\max e_2 \cdot \Pr(e_2 \text{ passes}) + 0 \cdot \Pr(e_2 \text{ fails})$$
$$= e_2 \cdot \Pr(F(e_2|y) < 1/2) \ .$$

The probability here is given by the setter's posterior over y, which is updated from his initial prior, $y \sim g(\cdot)$, by using information obtained from the first vote. How this posterior is formed must now be specified.

In the second period, the agents will have observed the outcome of the first referendum, specifically

$$V_1 = \text{proportion of the vote in favor of } e_1 \ .$$

Agents will use this observation to update their priors over the parameters of the taste distribution, y, using Bayes Rule. The vote in favor of e_1 was generated by some equilibrium voting strategy in the first period

$$v_1(e_1, t) = \begin{cases} 1 & \text{if } t \text{ votes in favor of } e_1 \\ 0 & \text{otherwise} \end{cases}$$

(which may be considerably more complex than the sincere voting strategy of the second period) and the actual realization of the parameter, y according to

$$V_1 = \int v_1(e_1, t) f(t|y) \mathrm{d}t \ .$$

Agents use this to update their priors over y after observing V_1. According to Bayes Rule, their posterior for y, $h(y|V_1, e_1)$, is given by

$$h(y|V_1, e_1) = \begin{cases} \dfrac{g(y)}{\int\limits_{A(V_1, e_1)} g(y)\mathrm{d}y} & \text{if } y \in A(V_1, e_1) \\ 0 & \text{if } y \notin A(V_1, \notin_1) \end{cases},$$

where $A(V_1, e_1) = \{y : V_1 = \int v_1(e_1, t) f(t|y)\mathrm{d}t\}$. Thus, updating in this environment takes the form of truncation of the prior to the set of y's, $A(V_1, e_1)$, that could have generated that vote in the first period[5].

Therefore, if the proposal was not approved in the first period, the setter will use this posterior over y to choose the second proposal, e_2. The setter will choose e_2 to

[5] This specification of Bayes Rule has implicitly assumed that the true distribution of tastes in the population is unchanged from one period to the next. In principle it could be generalized to allow for different, but correlated, distributions of tastes in each period. Different distributions are not modeled here because this is likely to be due to differences in voter turnout which in turn may depend on the proposals. However, Palfrey and Rosenthal's [10] model of rational turnout suggests that the distribution of voting citizens does not vary with the proposal. In any case, this is a model of foresighted voting rather than turnout.

solve

$$\max_{e_2} e_2 \cdot \Pr(e_2 \text{ passes} \mid V_1, e_1)$$
$$= e_2 \cdot \Pr(F(e_2|y) \leqq 1/2 \mid V_1, e_1)$$
$$= e_2 \cdot \int_{\{y: F(e_2|y) \leqq 1/2\}} h(y|V_1, e_1)\mathrm{d}y$$

Let $e_2(V_1, e_1)$ maximize this, and let its maximum value be $U^s(V_1, e_1)$, the expected value to the setter of what will occur in the last period if the proposal e_1 received a vote V_1 in the first period.

Having specified the possible outcomes for the last period given whatever may have occurred in the first period, it is now possible to define a first period equilibrium. The voters will observe the first period proposal, e_1, and vote according to whether they prefer this proposal or the expected value to them of whatever will occur in the future if the proposal fails. Given this voting stretegy for an arbitrary e_1, the setter will choose e_1 to maximize his current expected utility for a proposal plus the expected value to him of whatever will occur in the future.

The precise specification of the voters' first period strategy is more subtle than in the last period since there are more states to compare: the current e_1, the uncertain e_2 and reversion. As in the last period, a voter will vote for e_1 according to whether e_1 is preferred to its alternative, behaving as if the unlikely event occurred that the voter is decisive in the first referendum. In this case, however, the alternative is effectively a lottery between e_2 and 0 whose probability is the *ex post* probability that e_2 passes. The value to a voter of type t of this lottery is $u(e_2(V_1, e_1), t) \cdot \Pr(e_2 \text{ passes} \mid V_1, e_1)$, so a voter will compare this to $u(e_1, t)$ in deciding how to vote. Note, however, that the former depends on V_1, the fraction of the first vote in favor of e_1, which is unobserved at the time voter t must choose how to vote. By analogy with myopic voting in which a voter can affect the outcome only in the marginal case in which he is decisive, he should evaluate the lottery in this marginal case as well, $V_1 = 1/2$ (as opposed to the average or expected value of lottery valuations over either all possible V_1 or $V_1 < 1/2$, when e_1 fails).

This is expressed formally as follows: a voter in the first period will vote for e_1, $v_1(e_1, t) = 1$, according to

$$v_1(e_1, t) = \begin{cases} 1 & \text{if } \Delta(e_1, t) > 0 \\ 0 & \text{otherwise} \end{cases},$$

where

$$\Delta(e_1, t) = u(e_1, t) - u(e_2(1/2, e_1), t) \cdot \Pr(e_2 \text{ passes} \mid 1/2, e_1)$$

and $e_2(\cdot)$ and the *ex post* probability of passage are conditioned on $V_1 = 1/2$.

In fact, there are two ways in which a voter could, in principle, affect the outcome. One is the (negligible) case in which a voter is decisive and this is what is captured in Δ. Another function of a vote is in contributing to the aggregate signal, V_1, observed by the setter and used to choose e_2. That is, an individual vote contributes marginally to V_1 and will have a small effect in the (non-negligible) case in which e_1 fails, $V_1 < 1/2$. This signaling aspect is ignored in Δ primarily for reasons

of tractability. In its defense, however, it may be that the marginal case of being decisive is more important to a voter than the marginal signaling effects of an individual vote, in a sense which can be made precise. This will be true in the example below in which e_2 is not a function of V_1. Whether this is true more generally is difficult to establish with infinitely many voters. The following digression sketches a decision problem when there are finitely many voters and some chance that an individual is decisive and shows more precisely the decisive and signaling functions of a vote (and also shows more precisely why Δ captures the decisive function of a vote).

Maintain the other assumptions in the text, but suppose that there are $n+1$ voters, n even. Since there are a finite number of voters there is some small chance that a voter will be decisive and that a voting strategy can be strictly optimal in a voter's decision problem if there is no cost of voting, as is assumed. This is a Nash-type equilibrium in a voting game among the voters so this voter will take as given the voting strategies of all the other agents, and these can be summarized as

$$P_{V_1} = \text{probability that the other voters' first period vote will generate a fraction } V_1 \text{ in favor of } e_1$$

for $V_1 = 0, 1/n, 2/n, \ldots, 1$. These are generated by order statistics of the other agents' strategies conditional on tastes, and will generally not be zero since this voter does not know the tastes of the others. Of special interest is $P_{1/2}$, the probability that the other voters are evenly split for and against e_1 and this voter is decisive.

Let

$$U^t(V_1, e_1) \equiv u(e_2(V_1, e_1), t) \cdot \Pr(e_2 \text{ passes} \mid V_1, e_1)$$

$$= \text{a } t\text{-voter's } \textit{ex post} \text{ evaluation of the lottery the setter will choose if he observes } V_1 \text{ in favor of } e_1$$

$$v = \begin{cases} 1 & \text{if } t \text{ votes in favor of } e_1 \\ 0 & \text{otherwise} . \end{cases}$$

The voter must consider cases in which e_1 will pass irrespective of his vote ($V_1 > 1/2$), in which he is decisive ($V_1 = 1/2$) and in which e_1 will fail ($V_1 < 1/2$). If e_1 fails, the voter must also consider how his vote will add or detract from the aggregate percentage $(nV_1 \pm 1)/(n+1)$ which influences the future lottery. Taking P_{V_1} as given, the decision problem is to choose $0 \leqq v \leqq 1$ to

$$\max \sum_{V_1 > 1/2} P_{V_1} u(e_1, t) + P_{1/2} \{vu(e_1, t) + (1-v) U^t(n/2(n+1), e_1)\}$$

$$+ \sum_{V_1 < 1/2} P_{V_1} \{vU^t((nV_1+1)/(n+1), e_1)\} + (1-v) U^t(nV_1/(n+1), e_1)\} .$$

The solution to this problem is for the voter to vote for or against e_1 according to whether

$$P_{1/2} \{u(e_1, t) - U^t(n/2(n+1), e_1\}$$

$$+ \sum_{V_1 < 1/2} P_{V_1} \{U^t((nV_1+1)/(n+1), e_1) - U^t(nV_1/(n+1), e_1)\}$$

is positive or negative respectively. The first term embodies the decisive aspect of the agent's vote and the second term describes the signaling aspect.

If the signaling aspect is ignored or if, as in the example below the voters' *ex post* valuation does not depend on the first vote, then, if there is some chance a voter is decisive, as is likely to be the case with finitely many voters, the strictly optimal voting strategy is

$$v=\begin{cases}1 & \text{if} \quad u(e_1,t)-U^t(n/2(n+1),e_1)\geqq 0\\ 0 & \text{otherwise}\end{cases}.$$

As the number of voters increases, the probability of being decisive goes to zero, yet the limit of this voting rule is specified in Δ.

Whether the signaling aspect should be ignored depends on the behavior of the summation term in the expression above relative to $P_{1/2}$ as n increases. If this term goes to zero faster than $P_{1/2}$, then the signaling aspect is overwhelmed by the decisive aspect of a vote in large populations. This is assumed here. If the converse is true, then a voter will vote purely for its signaling value. Such an equilibrium can be considered objectionable on the methodological grounds that it is not desired to model voting in which an agent's preferences for the current proposal have *no* effect on strategy. The final case is one in which the ratio of the second term to $P_{1/2}$ converges to some non-zero constant, which would reflect a tradeoff between signaling and decisiveness. In view of the fact that the functional form of U^t (as well as its first argument) and P_{V_1} depend on n and t, so that each of these cases must be considered for every t, it is clear that the resolution of this issue must await an explicit analysis of finite populations[6].

Returning now to the specification of the equilibrium, it will be assumed that a voter considers the signaling aspect of his vote small relative to the decisive aspect[7].

Finally, given the first period voting strategy, the setter's foresighted choice of first proposal takes account of the (potential) trade off between winning approval and gaining information for the next round. It is specified last in the following summary.

Definition. A perfect Bayes equilibrium of this strategic voting game with learning consists of first and second period proposals, $e_1, e_2(V_1, e_1)$, first and second period voting strategies, $v_1(e_1, t)$, $v_2(e_2, t)$, and an *ex post* distribution of y, $h(y|V_1, e_1)$, such that

a) $h(\cdot)$ satisfies Bayes Rule:

$$h(y|V_1,e_1)=\begin{cases}\dfrac{g(y)}{\int\limits_{A(V_1,e_1)} g(y)dy} & \text{if} \quad y\in A(V_1,e_1)\\ 0 & \text{otherwise}\end{cases},$$

---

[6] This discussion also suggests that a fully strategic voter in the previous case of perfect learning may do more than simply compare the current proposal to the second period lottery. A careful analysis along the lines suggested here might even show that perfect learning is not possible.

[7] Morton [8] considers the tradeoff between signaling for future use versus current allocation in a two agent, private information bargaining model with a similar equilbirum concept.

where

$$A(V_1, e_1) = \{y : V_1 = \int v_1(e_1, t) f(t|y) \mathrm{d}t\};$$

b) $v_2(e_1, t) = \begin{cases} 1 & \text{if } u(e_2, t) \geqq u(0, t); \\ 0 & \text{otherwise} \end{cases}$

c) $e_2(V_1, e_1)$ solves for every V_1, e_1

$$U^s(V_1, e_1) \equiv \max_{e_2} e_2 \cdot \Pr(e_2 \text{ passes } | V_1, e_1)$$

where the probability of passage is calculated from $h(y)$, the *ex post* distribution;

d) $v_1(e_1, t) = \begin{cases} 1 & \text{if } \Delta(e_1, t) \geqq 0 \\ 0 & \text{otherwise} , \end{cases}$

where $\Delta(e_1, t) \equiv u(e_1, t) - u(e_2(1/2, e_1), t) \Pr(e_2 \text{ passes } | 1/2, e_1)$ and $e_2(\cdot)$ and the *ex post* probability of passage presume $V_1 = 1/2$; and

e) e_1 solves

$$\max_{e_1} E_y \begin{cases} e_1 & \text{if } \int v_1(e_1, t) f(t|y) \mathrm{d}t \geqq 1/2 \\ U^s(\int v_1(e_1, t) f(t|y) \mathrm{d}t) & \text{otherwise} , \end{cases}$$

where the expectation is with respect to the prior distribution $g(y)$.

5. The Construction of an Equilibrium

In this section, an equilibrium is constructed and explicitly characterized under some specific distributional assumptions. The general issue of existence and uniqueness is challenging, but this example should provide some insight. The assumptions are chosen to make the equilibrium easier to construct. The key simplification generated is that the second proposal will be deterministic in that it does not depend on the fraction of the vote in favor of the first proposal. That is, any distribution of tastes, y, which would vote down the first proposal would lead the setter to choose the same second proposal. (As remarked below, this implies that agents need only observe whether the first proposal failed, and not the precise fraction of the vote against.) This deterministic feature is likely not robust.

In this example, it will be assumed that it is common knowledge that voters' tastes are distributed normally[8] but the mean and standard deviation are unknown. Thus, the population parameters, y, are μ and σ, and $f(t|\mu, \sigma)$ is a normal density function. Further, all agents' priors, $g(y)$, over the population parameters are that μ and σ are jointly uniformly distributed over the rectangle $[0, m] \times [0, s]$, with $s = \infty$; i.e., diffuse.

This diffuse prior will remove a troublesome corner in the setter's second period problem and make the second proposal deterministic. A diffuse prior is objection-

[8] A negative t-type voter can be interpreted as a voter whose utility is even steeper than a $t = 0$ voter at $e = 0$. The upshot is that there is always a mass point of voters who prefer reversion to any positive expenditure level.

able since it does not allow updating via Bayes' Rule; also it is implausible that agents have this very special form of no beliefs about the parameter. Hopefully its use here will be considered less objectionable since Bayes voting strategies with finite but large support will be constructed, and it will then be argued that the limit of these is the appropriate strategy for the diffuse case.

A. The Form of Equilibrium Voting Strategies

To solve for this equilibrium a solution form for the voting strategies is conjectured. The second round is unproblematic: since there is no future, voters myopically vote their preferences for the second proposal versus reversion

$$v_2(e_2, t) = \begin{cases} 1 & \text{if} \quad u(e_2, t) \geqq u(0, t) \quad \text{or} \quad t \geqq e_2 \\ 0 & \text{otherwise} \ . \end{cases}$$

For a fixed e_2, this voting strategy is monotonically increasing in t.

The first round of voting is more problematic since voters' preference comparisons are more complex. It is

$$v_1(e_1, t) = \begin{cases} 1 & \text{if} \quad \Delta(e_1, t) \geqq 0 \\ 0 & \text{otherwise} \ , \end{cases}$$

where

$$\begin{aligned} \Delta(e_1, t) &\equiv u(e_1, t) - u(e_2(1/2, e_1), t) \cdot \Pr(e_2 \text{ passes } | 1/2, e_1) \\ &= (e_1 - e_2 \Pr)\, t - (e_1^2 - e_2^2 \Pr) \ . \end{aligned}$$

The first voting strategy is monotonically increasing (decreasing) in t if the coefficient on t in Δ is positive (negative). It will be conjectured in this section that it is increasing and then verified that the setter will choose a first proposal greater than the *ex post* probability weighted second proposal. Given this conjecture, for any first proposal, e_1, there is some switching type, $t^*(e_1)$, such that

$$v_1(e_1, t) = \begin{cases} 1 & \text{if} \quad t \geqq t^*(e_1) \\ 0 & \text{otherwise} \ . \end{cases}$$

The switching type, $t^*(e_1)$, will be the type of the decisive or indifferent voter, for whom $\Delta(e_1, t^*(e_1)) \equiv 0$.

The implications of these voting strategies, including the form of Bayesian updating, the setter's choice of proposals and the determination of $t^*(e_1)$, can now be worked out.

B. Bayesian Updating

All agents learn in the same way from the first round of voting since the information structure is common knowledge. Agents observe the percentage of the vote in favor of the first proposal and using the first period equilibrium voting strategy gain information about the unknown distribution of tastes in the population. The vote in

favor is determined by

$$V_1 = \int v_1(e_1, t) f(t|\mu, \sigma) \mathrm{d}t$$
$$= 1 - F(t^*(e_1)|\mu, \sigma)$$

since it is conjectured that voters with a type greater than $t^*(e_1)$ vote in favor. The agents can calculate the set $A(V_1, e_1)$ of all combinations of μ and σ which could have generated the vote V_1.

$$A(V_1, e_1) = \{\mu, \sigma \in [0, m] \times [0, s]: V_1 = 1 - F(t^*(e_1)|\mu, \sigma)\} .$$

The set $A(V_1, e_1)$, which is the support of the posterior, takes a simple form since tastes are distributed normally. Let ϕ be the standard normal cumulative and let τ satisfy

$$V_1 = 1 - \phi(\tau) .$$

Note that if $V_1 < 1/2$ (e_1 did not pass) then $\tau > 0$. Then μ, σ could have generated V_1 only if

$$\frac{t^*(e_1) - \mu}{\sigma} = \tau$$

and so

$$A(V_1, e_1) = \{\mu, \sigma \in [0, m] \times [0, s]: \mu = t^*(e_1) - \tau\sigma\} .$$

Bayesian updating in this model is simply truncation of the prior support $[0, m] \times [0, s]$ to that part of line segment $\mu = t^*(e_1) - \tau\sigma$ lying in the rectangle (see Fig. 2). Since the prior is uniform, so also is the posterior over the line $A(V_1, e_1)$. The slope, $-\tau$, is determined by V_1 and is negative if e_1 did not pass. The slope varies from the horizontal (if $V_1 = 1/2$) to the vertical (if $V_1 = 0$). It is the normality of the distribution of population types that yields the linearity of the posterior support of population parameters.

It will be assumed in this section that the equilibrium is such that $t^*(e_1) < m$, as shown in Fig. 2, and verified that such an equilibrium exists. The next section will indicate that the equilibrium is unique.

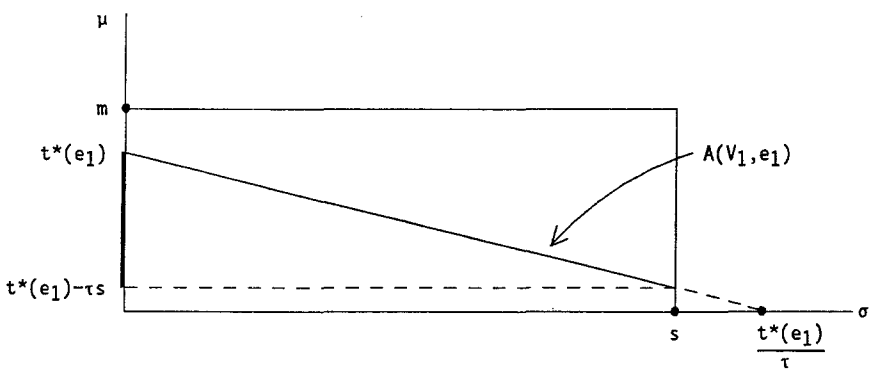

Fig. 2. Prior and posterior support of (μ, σ)

C. Second Period Equilibrium

In the second period, the setter will choose a second proposal optimally, given the posterior resulting from the first period. Whatever may have occurred in the first period, the voting strategies for a final proposal are sincere: a voter of type t will vote for e_2 if he prefers it to the reversion level of zero. According to the preferences assumed here

$$v_2(e_2, t) = \begin{cases} 1 & \text{if } t > e_2 \\ 0 & \text{otherwise} \end{cases} .$$

Consequently, e_2 will pass iff more voters are above e_2 than below, or the median of voter tastes is greater than e_2. The setter is still unsure of the true population characteristics, so, using the posterior, he will choose e_2 to

$$\max e_2 \cdot \Pr(\text{median}\{t\} > e_2 | A(V_1, e_1))^9 .$$

That tastes are distributed normally once again provides considerable simplification, since for the normal, the median is equal to the mean; so the setter's problem is

$$\max e_2 \cdot \Pr(\mu > e_2 | A(V_1, e_1)) .$$

The setter cares only about the μ-marginal distribution of the posterior. His uncertainty about σ serves only to prevent him from gaining perfect information about μ from the first period vote. In Fig. 2, the μ-marginal is uniform on the projection of $A(V_1, e_1)$ onto the μ axis. The highest possible μ is just $t^*(e_1)$ and the lowest depends on τ and is $\max\{0, t^*(e_1) - \tau s\}$. The μ-marginal is uniform over this support since the joint posterior is. The setter's problem is choose $e_2 \geqq \max\{0, t^*(e_1) - \tau s\}$ to

$$\max e_2 \cdot \frac{t^*(e_1) - e_2}{t^*(e_1) - \max\{0, t^*(e_1) - \tau s\}} .$$

The solution to this problem is straighforward. The setter will choose $e_2 = t^*(e_1)/2$ unless this is less than the lowest proposal that is sure to pass, in which case he chooses that proposal $e_2 = t^*(e_1) - \tau s$. The solution is summarized below:

	e_2	$\Pr(e_2 \text{ passes } \mid A(V_1, e_1))$
$0 < \tau < t^*(e_1)/2s$	$t^*(e_1) - \tau s$	1
$t^*(e_1)/2s < \tau < t^*(e_1)/s$	$t^*(e_1)/2$	$t^*(e_1)/2\tau s$
$t/(e_1)/s < \tau$	$t^*(e_1)/2$	1/2

It is convenient to restrict attention to the last case since e_2 and its *ex post* probability of passage are not functions of the first period vote (or τ) and so are not random from the point of view prior to the first vote. The reason this is convenient is

[9] Since updating here is simply truncation, posterior probabilities will hereafter be written conditionally on the set $A(V, e_1)$, which is the common information set (in the sense of extensive games) at the beginning of the second round.

that the next step in constructing the equilibrium is to solve for the cutoff voter $t^*(e_1)$, who is indifferent between e_1 and the expected value of the second period and this step is considerably easier if the *expected* value is not a function of the first vote. This simplification is made by assuming that agents behave according to the above decision rule with $s=\infty$, i.e., as if only the third case obtains, in which $e_2=t^*(e_1)/2$ and Pr (e_2 passes $|A(V_1, e_1))=1/2$, independent of τ (or V_1). It is not possible to solve for the equilibrium directly when $s=\infty$ since Bayesian updating is not strictly defined when the prior is diffuse. This decision rule is chosen since it is the limit of the more general one as s grows large[10].

This nonstandard prior for σ is used here for the sake of tractability of the resulting decision rules. If the use, or this approach to "updating", a nonstandard prior is objected to, note that the same decision rules can be justified analytically by assuming that agents observe only whether the first proposal passed or failed, rather than the size of the vote in favor[11].

To see this, suppose it is only observed that $V_1<1/2$; i.e., that the first proposal failed. Then all that can be inferred about τ (defined above) is that $\tau<0$. Then the support of the posterior, which was the line $A(V_1, e_1)$ when V_1 and τ were known, is now the subrectangle $[0, t^*(e_1)]\times[0, s]$. The posterior is uniform over this subrectangle and the μ-marginal is uniform over $[0, t^*(e_1)]$. So the solution to the setter's second period problem is the third case above, as was desired. Consequently, in this environment, the decisions of agents with a diffuse prior over τ who observe the first vote are equivalent to those with a standard prior who observe only success or failure of the first proposal.

D. First Period Voting Strategy: Determination of the Switching Type

The second proposal and *ex post* probability of passage have been determined as a function of the switching voter type, $t^*(e_1)$. Now first round voting can be examined as a function of e_1. Voters favor the first proposal if they prefer it to the *ex post* expected utility of the second proposal, which is a function of the switching type. But from $\Delta(e_1, t^*(e_1))\equiv 0$, the switching type, in turn, is indifferent between obtaining e_1 now or his ideal point $e_2=t^*(e_1)/2$ with *ex post* probability 1/2 later. Therefore, $t^*(e_1)$ is given by $\Delta(e_1, t^*(e_1))\equiv 0$ which becomes

$$u(e_1, t^*(e_1))-u(e_2, t^*(e_1))/2=0$$

or

$$(e_1 t^*(e_1)-e_1^2)-(t^*(e_1)^2/2-t^*(e_1)^2/4)/2=0 \ .$$

The solutions to this quadratic equation are

$$t^*(e_1)=2e_1(2\pm\sqrt{2}) \ .$$

---

[10] The limit here is analytical, not a probability limit: for any finite s, the *ex ante* probability that one of the first two cases obtains, given that e_1 fails, is equal to 1/2, so these cases never become negligible. It is also worth noting that although it has been argued that a voter will be concerned with only those cases in which he is decisive – in which the vote in favor is just less than 1/2 and the first case obtains – nevertheless if $s=\infty$, then $A(V_1, e_1)$ will eventually cross the horizontal axis for any $V_1<1/2$.

[11] Thanks to a referee for suggesting this.

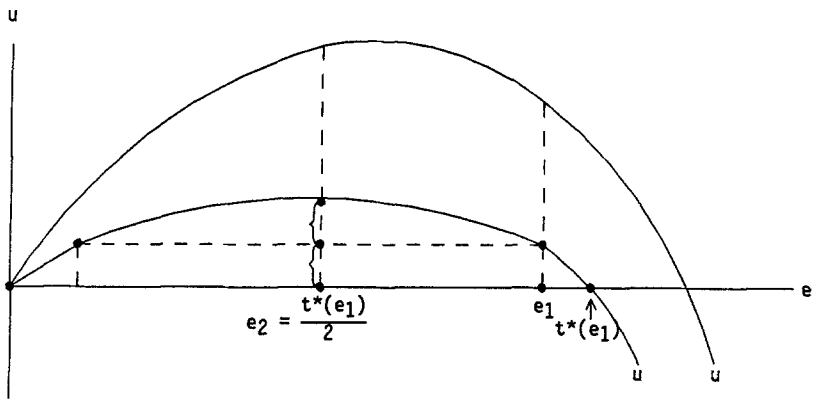

Fig. 3. Determination of switching type

One of these solutions is ruled out by the additional consistency criterion that voters above $t^*(e_1)$ strictly prefer e_1; i.e., that the coefficient on t in Δ be positive, or $e_1 > e_2/2$. Simple substitution implies that the only consistent solution is

$$t^*(e_1) = 2e_1(2 - \sqrt{2}) \ .$$

The solution is pictured in Fig. 3. Although causuality runs in the opposite direction is is easiest to understand by fixing $t^*(e_1)$ and seeing what e_1 would have to be to support that switching type. If $t^*(e_1)$ is the cutoff, then the setter will find his *ex post* utility greatest by choosing $e_2 = t^*(e_1)/2$, the cutoff's ideal point. However e_2 will pass will with only a 50% probability, so his expected utility is only half that of his ideal point. The certain utility value to $t^*(e_1)$ must be equal to this; there are two values of e_1 which do this, one above and one below e_2. However the less of these is ruled out by the additional condition that types above $t^*(e_1)$ strictly prefer e_1.

Note that for any $e_1 > 0$, $t^*(e_1) > e_1 > e_2$. Like Romer and Rosenthal's naive equilibrium, the proposals decrease over time. (The next section considers whether there is an equilibrium in which they increase.) What distinguishes the sophisticated equilibrium here is that the switching type is strictly greater than the first proposal – types in $[e_1, t^*(e_1)]$ prefer e_1 to reversion but will nevertheless vote against e_1 in anticipation of a better second proposal: the sophistication of voters limits the power of the setter's agenda control.

E. The First Proposal

To fully characterize the equilibrium, it remains only to find the setter's optimal first proposal. The setter want to choose e_1 to maximize his expected utility of the first period plus his expectation of the (ex post) expected utility of what will occur in the second period if e_1 fails. Specifically, choose e_1 to

$$\max e_1 \cdot \Pr(e_1 \text{ passes}) + (1 - \Pr(e_1 \text{ passes})) \cdot e_2 \cdot \Pr(e_2 \text{ passes} \mid \Delta(V_1, e_1)) \ ,$$

where e_2 is a deterministic function of e_1 given the first period voting strategy and

$\Pr(e_2 \text{ passes} \mid A(V_1, e_1))$ is evaluated by the setter's posterior conditional upon e_1 failing; on the other hand, $\Pr(e_1 \text{ passes})$ is evaluated at the setter's prior.

Since the equilibrium voting strategy is increasing in t, e_1 will pass iff the greater mass of voter is above $t^*(e_1)$ or median $\{t\} > t^*(e_1)$. As before, tastes are distributed normally and the median is equal to the mean so e_1 will pass if $\mu > t^*(e_1)$. Once again, the setter cares only about the mean, and the μ-marginal of his prior is uniform over $[0, m]$, so $\Pr(e_1 \text{ passes}) = \Pr(\mu > t^*(e_1)) = 1 - t^*(e_1))/m$. The setter's problem becomes

$$\max e_1 \cdot (1 - t^*(e_1)/m) + t^*(e_1) e_2/2m$$

$$= e_1 \cdot (1 - e_1 2(2 - \sqrt{2})/m) + e_1 2(2 - \sqrt{2}) e_1 (2 - \sqrt{2})/2m$$

and the solution is

$$e_1 = m/4(\sqrt{2} - 1) \ .$$

The *ex ante* probability that e_1 passes is

$$1 - t^*(e_1)/m = 1 - 1/\sqrt{2}$$

The setter's *ex ante* expected utility is

$$m(\sqrt{2} + 1)/8 \ .$$

This is greater than $m/4$, the setter's maximum expected utility, if there was only one round of voting. This result is useful in proving uniqueness of equilibrium.

Having found the "sophisticated" voting equilibrium, it is possible to make some comparisons with the "naive" case. Romer and Rosenthal anylzed repeated referenda under the assumption that *all* agents behaved myopically; that is, voters only compare the current proposal to reversion, ignoring any future proposals, and *all* agents, including the setter, do not revise their priors in the light of previous votes. Let x_1, x_2 be the proposals for this "naive" case. In each stage x_i will pass if $x_i < \mu$ so the *ex ante* probability of passage is

$$\Pr(x_i \text{ passes}) = \Pr(x_i < \mu) = (m - x_i)/m \ .$$

The setter then chooses the proposals x_1, x_2 to

$$\max x_1 \Pr(x_1 \text{ passes}) + (1 - \Pr(x_1 \text{ passes})) x_2 \Pr(x_2 \text{ passes})$$

$$= x_1 \frac{m - x_1}{m} + \frac{x_1}{m} x_2 \frac{m - x_2}{m}$$

where all the probabilities are evaluated according to the prior. The solution is

$$x_2 = m/2 \ ; \quad x_1 = 5m/8 \ .$$

So, in comparison,

$$x_1 > e_1 > x_2 > e_2 \ ,$$

"naive" behavior increases the proposal sequence. The differences between the two formulations of the bureaucrat's decision problem are that e_2 is based on learned or

ex post information whereas x_2 ignores any information implicit in the first vote (x_2 is the proposal that would be chosen if the setter had only one chance to obtain approval); and e_2 must win approval from strategic voters and passes only if the median type is greater than $t^*(e_1) > e_1$, whereas x_1 faces naive voters and will pass if the median type is greater than x_1.

6. Uniqueness of Equilibrium

The strategic equilibrium constructed in the last section is unique and this is the subject of this section. The proof consists of a tedious consideration of cases, the details of which appear not to offer much insight, so the proof is only briefly sketched here[12]. Three main cases are outlined; each either violates a consistency condition or can be prevented by the setter who does better under the equilibrium of the last section.

One structural feature of the model does seem to be robust however: the monotonicity of the first period voting strategy. Voters vote for or against e_1 according to whether Δ is positive or negative. Since Δ is monotone (in fact, linear for quadratic utility) in t, this suggests that there will be a unique cutoff voter type, $t^*(e_1)$, who is indifferent between e_1 and whatever may occur next, defined by the type who makes Δ equal to zero. Types on each side of the cutoff will vote together, for or against e_1, according to whether the coefficient on t in Δ positive or negative. This considerably simplifies the analysis since the aggregate vote in favor of e_1, which the setter used to obtain information about the population in choosing the second proposal, will then be either the mass of voter types above or below the cutoff type.

This also suggests that there may be two equilibria in this model, according to whether the coefficient on t in Δ is positive or negative: roughly, whether e_2, weighted by the *ex post* probability of passage, is less than or greater than e_1. Of course, equilibrium voting strategies must satisfy an additional consistency condition, that the information gleaned by the setter from observing the first period vote will lead him to optimally choose an e_2 which provides the correct sign on the coefficient of t: it must be positive it voters above the cutoff vote for e_1, and negative if voters above the cutoff vote against e_1. If this were true, then, bearing in mind that the cutoff voter is indifferent, it might be said that voters above the cutoff voted for (or against) the first proposal because they expected the value to them of the second proposal to be lower (or higher) than the first; further, use of this strategy to understand how a first period vote was generated should lead the setter to choose a second proposal which was, in fact, lower (or higher) than the first. The second type of equilibrium fails under the normality assumptions made here, but whether it fails more generally is an open question. The first equilibrium may seem more natural since it is qualitatively similar to the "naive" voting strategies found by Romer and Rosenthal for multiple referenda. However, the second type is not at all analytically unreasonable – voters who prefer big budgets vote against the first proposal

[12] Full details are available from the author.

precisely because they expect the second to be higher – and may be possible under different assumptions about the distribution of voter types.

Before considering the case of strictly decreasing voting strategies, in which the coefficient on t in Δ is negative, it is useful to first examine the contrary of the assumption made in the last section that the switching type, $t^*(e_1)$, is less than m, the highest possible mean of the population. If $t^*(e_1) \geqq m$, then the posterior support line $A(V_1, e_1)$ cuts across the top ($\mu = m$) of the prior support rectangle instead of the left edge ($\tau = 0$) as in Fig. 2. The first proposal will surely fail $\mu \leqq m \leqq t^*(e_1)$. Further assuming s is infinite the posterior will eliminate some of the possible τ's but none of the possible μ's; so no new relevant information is gained from the first vote and the second period is equivalent to Romer and Rosenthal's second period. However the setter does better in the equilibrium of the last section and can force it by proposing that e_1. That is, assuming $e_1 = m/4(\sqrt{2} - 1)$ and no learning for the second period (and the resulting second proposal) implies $t^*(e_1) \approx 0.67m < m$.

Now consider the case in which the coefficient on t in Δ is strictly negative, so that voters *below* the switching type vote *for* the first proposal. A parallel analysis to that of the last section can be performed with the major difference that the posterior support line $A(V_1, e_1)$ is *upward* sloping so that information from a failed first proposal truncates the set of possible μ's from below. As before, cases in which $t^*(e_1) \gtrless m$ need to be considered but also whether $t^*(e_1) \gtrless m/2$ must be distinguished. This last corresponds to whether there is a corner solution to the setter's second period problem which is a possibility if the posterior distribution on μ has been truncated away from zero. Although tedious, it is straightforward to show that in each case either a) a consistency condition is violated in which the setter's optimal choice of e_1, e_2 and the *ex post* probability of passage violate the initial assumption that the coefficient on t in Δ is negative; or, although consistent, the setter can enforce the equilibrium of the last section in which he does better.

Finally, for completeness, the case in which the coefficient on t in Δ is identically zero can be considered. Then voting strategies do not depend on t and all voters vote for, against, or entirely randomly depending on whether the intercept term in Δ is positive, negative or zero. In each case, no new information is communicated from the first period vote so the second proposal and *ex post* probability are determined myopically, and the first proposal is determined since the coefficient is zero. Based on these, the intercept term is positive and the first proposal passes for sure. But once again, the setter can enforce the equilibrium of the last section in which he does better.

7. Conclusion

This model has begun to address issues in the strategic acquisition and use of information in a voting process using a perfect Bayes equilibrium concept. Although the simplest framework might be thought to be a pure public opinion poll (or "straw vote") prior to a referendum, the results of Crawford and Sobel [2] suggest that voters would completely mask their preferences. Instead, a repeated referendum process is analyzed which requires voters to take the first proposal more seriously as a possible allocation, rather than a mere signal. This additional

structure generates an equilibrium in which learning does occur, and is at least quantitatively different from naive or myopic behavior. This concluding section will discuss the robustness of the equilibrium constructed and its empirical implications.

Equilibrium depends upon some notion of consistency (or "rational expectations") between the information on the population contained in the first vote as determined by first round voting strategies and the implied second proposal optimally chose by the setter given that information. What makes this connection relatively simple is the nature of Δ, a voter's comparison of the current proposal to the uncertain future. Because there is a continuum of heterogeneous and noncooperative voters, equilibrium voting strategies are pure and an equilibrium can be characterized by the cutoff voter type who is just indifferent in this comparison. This aspect of equilibrium is likely to prove robust to generalizations. For example, with other (single peaked and concave) utility functions over a one-dimensional issue space, Δ will not be linear, but should retain monotonicity; single peakedness is generally interpreted as agents being in agreement in preferring "moderate" to "extreme" outcomes so the comparison expressed by Δ should switch signs for different types at most once. Then an equilibrium can be characterized by the switching type and whether the derivative of Δ with respect to type is positive or negative. If all equilibria are of this form, as seems likely, then this will provide an important simplifying step in a more general analysis.

If all equilibria do take this simple form, then another feature of the analysis follows from this. Bayesian updating is simply truncation of the prior to the set of parameters, y, which could have generated the first vote V_1:

$$V_1 = 1 - F(t^*(e_1)|y)$$

or

$$V_1 = F(t^*(e_1)|y)$$

depending on whether Δ is increasing or decreasing in t. In the example, F was the normal cumulative and y its unknown mean and standard deviation, so the support of the posterior was a straight-line segment. Presumably, other population distributions can be plausibly parameterized so that the posterior support is very well behaved.

An advantage of the normality assumption in that example is that only one of the parameters, the mean, determined whether a proposal would pass at each stage. Uncertainty about the standard deviation served only to prevent agents from gaining perfect information after the first vote. Although it will not generally be true for other distributions and priors that only one of the parameters determines the outcome, it would seem that, in most cases, observation of a vote serves to reduce the parameter space by one dimension. Presumably, if y is an n-dimensional vector, agents would have perfect information (as in the case of perfect learning) after observing the results of $n-1$ referenda.

Thus, the relative simplicity of equilibria would appear to hold in a larger class of environments. On the other hand, explicit, closed form solutions may be harder to come by. Such solutions were obtained in the example here only by a very careful choice of priors.

A deep problem which had to be confronted, if not solved, is the explication of rational voting in a continuum when no voter has weight in the outcome. It was argued that two common heuristics – voting as if one were decisive and voting one's preferences – were not equivalent, and the decisive and signaling functions of a vote were distinguished. For reasons of tractability, the decisive function was emphasized in the definition of equilibrium, although in the normal example constructed, the second proposal was not a function of the first period vote, so that there was in any case no signaling function to capture. A careful limiting analysis of finite voting populations might show some trade off in these functions in the voter's problem. Yet the previous remarks on the monotonicity of Δ and the role of a unique switching type would still hold. It is possible, however, that the equilibrium could unravel.

The equilibrium found here is different from that implied by naive or myopic voting behavior; the strategic proposal sequence is uniformly lower than that found by Romer and Rosenthal. Using their market structure analogy, one might say that strategic behavior by the voters reduces the "monopoly" power of the setter. However, there is not any apparent qualitative difference between the two equilibrium behaviors. To qualitatively distinguish them, one must consider the possibility that there may be another strategic equilibrium in which voting strategies are decreasing in t. Although the proof that this is not possible for the example considered was sketched in the last section, it is likely not generalizable.

A voting strategy will be decreasing in t if voters who prefer low budgets vote for the first proposal because they expect the second to be higher. From the setter's point of view, the first proposal not only yields information but (unlike a pure public opinion poll) also may succeed. If low and high proposals are, in some sense, equally informative then the setter might prefer to experiment with a high proposal with a desirable outcome if it succeeds. On the other hand, one can imagine somewhat contrived (and therefore uncommon) information structures in which lower proposals generate more informative signals; in the extreme, proposals below some level may be perfectly informative. In such a case, the setter may determine that the value of acquiring good information will outweigh the risk of a low first proposal succeeding.

Decreasing voting strategies and increasing proposals sequences are uncommon but not unreasonable and data suggests that it is at least worth entertaining the possibility. Romer and Rosenthal report that in a small but significant number of cases of local school budget referenda in Oregon, later proposals were larger than earlier defeated proposals. In their model of naive voting, proposal sequences should always decrease. It would appear that the only hope for a rational explanation of increasing proposal sequences rests with strategic voting behavior. If voters are strategic in these rare cases, one might argue that they are more generally as well.

References

1. Cramton PC (1983) Bargaining with incomplete information: An infinitehorizon model with two-sided uncertainty. Rev Econ Stud 51:579–593
2. Crawford VP, Sobel J (1982) Strategic information transmission. Econometrica 50:1431–1451

3. Fudenberg D, Tirole J (1983) Sequential bargaining with incomplete information. Rev Econ Stud 50:221–247
4. Kreps D, Wilson R (1982a) Sequential equilibria. Econometrica 50:863–894
5. Kreps D, Wilson R (1982b) Reputation and imperfect information. J Econ Theory 27:253–279
6. Ledyard J (1984) The pure theory of two candidate elections. Public Choice 44:7–41
7. Milgrom P, Roberts J (1982) Predation, reputation and entry deterrence. J Econ Theory 27:280–312
8. Morton S (1984) A dynamic revelation analysis: trading for information. Mimeo
9. Myerson R (1981) Optimal auction design. Math Oper Res 6:58–73
10. Palfrey T, Rosenthal H (1985) Voter participation and strategic uncertainty. Am Polit Sci Rev 79:62–78
11. Romer T, Rosenthal H (1979) Bureaucrats versus voters: on the political economy of resource allocation by direct democracy. Quart J Econ: 563–587